Reflections on Wireless Sensing Systems: From Ecosystems to Human Systems

Deborah Estrin

1. Introduction

We have learned a lot in the past decade and many of the issues we thought would be critical have turned out to be less so; but on the other hand a host of challenges and design and application opportunities have arisen that we certainly did not anticipate, and definitely would not have discovered had we not launched into the field. And if we are lucky, the reality will be more important than the original target.

First off, ***why*** wireless sensing systems, or what we often call Embedded Networked Sensing? Miniaturization and Moore's law have enabled us to combine sensing, computation and wireless communication in integrated, low-power devices, and to embed networks of these devices in the physical world. By placing sensing devices up close to the physical phenomena we are now able to study details in space and time that were previously unobservable. Across a wide array of applications, the ability to observe physical processes with such high fidelity will allow us to create models, make predictions, and thereby manage our increasingly stressed physical world. Most generally stated, our objective is to maximize information return from these adaptive sensing and actuation systems, across design, deployment, and run time; and more specifically through the design of multiscale, multimodal, and in-network processing algorithms. The essential power of this technology derives from **embedding** measurement devices in the physical world and **networking** them to achieve intelligent, adaptive, coordinated **sensing** systems. ENS has the perfect ingredients for multidisciplinary research because it offers transforming capabilities to the applications and challenging problems for the technologists. Most generally stated our objective is to maximize information return from these adaptive sensing and actuation systems, across design, deployment, and run time; and more specifically through the design of multiscale and in-network processing algorithms.

To my knowledge, UCLA's own bill kaiser and greg pottie designed the first integrated wireless sensor nodes about a decade ago under DARPA funding. Currently, the basic hardware and software building blocks to realize wireless sensing systems and enable associated applications is available. And while we will continue to see innovation in these fundamental components, we are at a very exciting juncture where we can now address what one of my Ecologist colleagues, Phil Rundel, refers to as ***2nd generation questions***. As an ecologist, he means that the observational capabilities of the technology are enabling the scientists to start to ask questions that they could not previously consider (not just answer prior questions better). In addition, from the technologist's perspective the technology designers can start to ask new technological questions that go beyond the basic system building blocks to designing the next generation of data intensive and adaptive systems.

2. Environmental observatories

While many of our initial conceptions of the field have seasoned and shifted over the years, environmental sensing applications and the growth of environmental observatories have proven to be a powerful application driver for wireless sensing systems**.** Many critical environmental processes exhibit high frequency spatial variation and are tremendously heterogeneous, requiring dense sensing in space and time. The spatial and temporal sampling density of ENS enables scientists to develop richer models that capture this heterogeneity and will support many applications whose models indicate the need for monitoring and management at spatial scales not achievable with remote or manual sensing.

We have focused on 4 science application domains--seismic, terrestrial ecology, aquatic, and contaminant transport--and in recent years have had some interesting real world (as in messy!) deployments. These science applications have taken us beyond the rhetoric to the even more exciting reality of ENS, and have also taken us across the globe to explore its applications. Following are some motivating examples:

- Our **contaminant transport** application has taken us to Bangladesh, where a computer science and an environmental engineering PhD student, developed and used CENS soil pylon technology and innovative system software and statistical tools to study the processes underlying the arsenic poisoning of drinking wells. In particular, how rice paddy irrigation may contribute to dissolved arsenic influx into aquifers. Just under 50 sensors were deployed over 12 days, collecting over 25,000 measurements. In addition to collecting some interesting scientific data, the system deployment gave us great insights into how fault detection should be designed for rapidly deployed sensor networks and how such systems can increase the quantity and quality of recovered data.
- A few months before this deployment, on October 28th, 2005, in nearby Pakistan, a major, 7.6, earthquake was detected by our **Seismic** array, that had recently been deployed in Mexico as part of a collaboration with Caltech and the Moore foundation. Our system of 50 wireless seismometers is the first array of its kind and has provided high spatial resolution observations; while driving development of disruption tolerant networking systems. Interactive interfaces operating over intermittently disconnected networks provide visibility into the long distance wireless links (~5km or more) and the supporting infrastructure. This is the primary requirement for almost every networked sensing system: the ability to know whether the system is up and running and collecting usable data, even if the analysis of that data is not done until a later date.

- Returning to California, NIMS (a robotic sensing node) and Multiscale actuated sensing technologies were applied to a study of the San Joaquin and Merced rivers, achieving previously unprecedented spatial and temporal observation of river mixing. Specific conductivity, ph, water velocity, and nitrate profiles that have not been previously observable are obtained with a portable system. The resulting high spatial resolution profiles can inform important policy regarding land use and water resources. This example illustrates the importance of rapid deployment technologies that are effectively mobile at multiple scales.
- Terrestrial and aquatic observing systems, which combine static wireless and multiscale actuated technologies, have been strategically deployed and evaluated at the James Reserve in Idylwild. Networks of imagers, above and belowground microclimate, as well as NIMS based imaging and aquatic systems are developed and employed to observe the environment. In addition to serving as one of our preliminary and longest running deployments, the data collected from these sites have served as a driver for rich data visualization development. Critical for analysis and even for guiding subsequent sensor placement. . Our terrestrial ecology instrumentation combines imaging and microclimate observations at multiple scales. From relatively broad field of view imagers on towers to nest-box imagers, these image sensors are one of our few deployable “biological” sensors and are motivating some very interesting computer science algorithm development to enable automatic filtering and event detection.
- Finally, our marine biologists are using data from one of the mobile aquatic applications that combines static buoy based sensors with a mobile robotic boat, affectionately referred to as roboduck. This system is enabling multiscale environmental and biotic observation of aquatic systems. In particular, the NAMOS deployment on the Lake Fulmore uses automated sample-collection from robotic nodes combined with static nodes to study phytoplankton dynamics.

3. Field inspired systems research

Through this intensely multidisciplinary research process we learned that a “measurement goal” is not sufficient; we must include the science question to identify the required precision, sampling regimen, synchronization. More importantly, the system must be sufficiently robust to these application requirements and the deployment environment to collect useful data (although it might well need babysitting by its developers).

Accompanying this focus on feedback from our applications is the challenging but rewarding process of learning from the field. We learned to let go of early assumptions and favorite ideas and enjoy what reality presented. For example Directed Diffusion, our initial notion of routing, was inspiring but in retrospect was probably too general and complex for initial generations of systems. However, the ideas of data centric algorithms, tasking, and tree based collection introduced in that paper are fundamental to the way these systems work. We often find ourselves in a tension between wanting to project forward so as to not work on problems that are made irrelevant with time, and yet to not end up exploring a space that is founded on speculation that might not come true.

We have learned much over the last few years of experimental systems development and deployment. We have learned that it is not just about enabling the largest number of smallest, lowest power, devices to create fully autonomous systems that are maximized for longevity, but rather it is about optimizing the end to end heterogeneous system. Nor, is it just about minimizing energy use and bits transmitted, but rather its about minimizing sensing uncertainty. Consequently, we have developed architectures and algorithms that optimize the system as a whole across a heterogeneous mix of (a) components, that include power on demand platforms, (b) sensing modalities that combine physical and imaging, and (c) mobile as well as static sensors. **In fact, mobility has turned out to be a key to overcoming the inherent undersampling of static sensors.** Where robots and automated sensors can not themselves maximize information return, we have begun to more explicitly design the systems to accommodate the human tier to enable data collection and observation that is guided by and guides other system resources. This has led us to focus on interactivity and in network processing that supports such system responsiveness. And a final pervasive and critical theme is data quality and system integrity for which in-network processing algorithms and powerful statistical techniques are proving to be the first steps to calibration, self test and validation. Let me take a few minutes to address a few of these themes, which happen to be of particular relevance to Mobicom.

Our approach to achieving scalable and robust systems is through designs that explicitly **exploit heterogeneity** and hierarchy. The smallest nodes, commonly referred to as motes, can be distributed most densely. However, they necessarily have limited sensing and computation, and communication range--Therefore for every flock of small nodes, we place microservers that have larger energy resources and correspondingly higher end sensors, processors, storage and communication. Going beyond this two- tier system of static nodes, one of the CENS innovations has been the broad introduction of a third tier in this hierarchy, that of autonomous mobile nodes. This tier of robotic capability allows adaptive sensing and sampling; in fact it is **only** with mobility that we have achieved **dense** sensing. Articulation greatly magnifies the effective sensor range, even when over small distances (such as a pan/tilt camera). And making use of infrastructure to assist the mobility makes it feasible in more extensive situations.

A second theme I want to emphasize is **Interactive** real time access to the embedded sensor systems and data in combination with other data sources. Instead of viewing ENS as a technology that is only useful when fully automated, we see it more like MRI technology, which when combined with image enhancing signal processing, other measurement modalities, and with human cognition, provides a tremendously powerful science and engineering tool. Moreover, it is through interactive access to these observing systems in the field, combined with innovative statistical tools, forming what my colleague Mark Hansen refers to as “coupled human-observational systems”, that we see rays of hope for addressing some of the most difficult issues of calibration and data integrity. As we saw from the mexico seismic array example, that one of the greatest values that wireless brings to sensors and instruments is the real-time **feedback** that they are working, or not, allowing for fast response and reduced data gaps, as well as opportunities for the user in the field to disambiguate confusing inputs.

Combining these themes of mobility to achieve density and user interactivity to achieve information return, is the increasing focus on **rapidly deployed (RD) systems**. RD systems are deployed for short durations, either repeatedly in the same location when periodic surveys are needed, or sequentially in different locations to achieve greater spatial extent. The presence of the user either continually or frequently facilitates timely calibration and maintenance as needed for some existing chemical sensors. It also enables triggered physical-sample collection, which when combined with laboratory analysis achieves the functional equivalent of dense bio/chemical sensing, where high density physical sensors record the precise local characteristics in which a physical sample is collected and then analyzed in the lab (wet chemistry and bio analysis). RD has become our most powerful usage model for environmental science and at the same time raises interesting technical challenges. In addition to system techniques that support rapid convergence and data extraction, these systems motivate the need for high system visibility, and statistical/design tools for incremental deployment and in situ data calibration and validation.

One of the commonalities across all of these themes of mobility and interaction and rapid deployment is the critical role that statistical techniques play in our systems. We see this in the problems of Experimental design and sensor layout where there are significant research opportunities to develop adaptive, iterative algorithms for deployment; however unlike earlier work this is focused much more on the sensor than wireless coverage. And as we scale up to large deployments in complex environmental media like soils, there is a need to apply geospatial statistical techniques for optimizing sensor placement subject to budget and regulatory constraints as well. As I said earlier, Data integrity is clearly one of the most challenging problems in sensor networks generally and there is much research to be done to develop robust procedures for aggregation and analysis. For example, as CENS folks have worked to interpret ENS collected data we have developed Spatio-temporal models in the form of flexible or “nonparametric” descriptions of signals such as for the urban creek nitrate runoff studies

Not unrelated to the increasing role for statistical techniques to turn noisy instruments and measurements into meaningful data, I want to mention a very interesting avenue for high impact research in the context of turning high sampling rate image and acoustic observations into interpretable measures (sensors)…This again came from field experience of realizing that available physical and chemical sensors in important cases just do meet the observational requirements of the application. To this end we started investigating imaging and acoustics as in situ biological sensors. For example, Mohammad Rahimi worked with Agilent to develop Cyclops, which is a 3 volt mote-based wireless imager (truly wireless in that it can work off of small batteries for extended periods, unike a webcam) and is intended to be programmed for local filtering and processing. The value of this approach is that there are many biological phenomena for which we do not and will not have direct measures for a long time, and yet which can be observed in the optical domain. And having easily deployed imagers offers the opportunity to combine observations from a diversity of poses, distances, angles, etc, and most importantly to address occlusions that are inherent to complex environments--in fact if you don’t have occlusions a single higher power imager can do better.

Of course there are tradeoffs in that these imagers return relatively low resolution images (for this day and age) and support relatively simple local processing. However, this is in fact part of what makes it interesting. In the target application domains its essential to make use of local context of the application and of the ambient environment to facilitate local processing and identification and so there is much opportunity for designing distributed algorithms that make use of higher end tiers in the system to adjust the local parameters

Similarly, moving from imaging to acoustics, we have the need for a hierarchical but distributed processing architecture (such as was suggested in the Govindan et al Tenet architecture) where the distributed devices are used for flexible in situ data collection, simple filters are applied to both conserve energy, reduce congestion, and adapt local sampling frequency. The microserver or master devices monitor observations and adapt local filtering parameters to match dynamic conditions and user requests. We have just begun to skim the surface when it comes to turning these high frequency data sources into “embedded sensors”

4. Participatory sensing

Our use of images and acoustics as sensing modalities started us thinking of the opportunity to use widely proliferated imagers and microphones in a similar manner, in particular, the imager-microphone-wireless-sensor packages that we all carry on our belts and in our pockets…the omnipresent cell-phone. It was this line of thinking along with the continued lesson that multiple scales of mobility is essential to achieving spatial density and extent, that started what we refer to as our participatory sensing research activities. Many of the tools, techniques, and research results that we have available to us to build upon in fact came out of the Mobicom community; as you well know.

While CENS continues to emphasize environmental sensing scientific applications, we have also begun to explore how these same design principles and techniques apply to sensing in the public sphere.And so far, Many of our architectural tenets and and lessons appear applicable to this context in which we hope to exploit the millions of sensors that are already carried and connected. In particular, in this context we are finding again that multiscale data and models are essential to design into the architecture to achieve context, and mobility is essential to achieve scalability and coverage. And interestingly enough in-network processing is essential to support privacy and personal control.

One of the most exciting aspects of this new focus for wireless sensing is the range of application types that we believe will bring sensing into every day life:. From "directed sensing" applications for self-administered health diagnostics, and participation in public and community health monitoring, to citizen sensing applications such as participatory urban infrastructure management, such as documentation of flooded storm drains, uneven sidewalks, potholes, poorly timed traffic lights, etc.

This range of applications is commonly enabled by the billions of cell phone users worldwide, combined with automated geocoding (GPS and cellular) and connectivity, the availability of image and acoustics as both data and metadata, the opportunity for local processing akin to what I just described for cyclops, and the ever increasing usability and sophistication of geospatial interfaces for data authoring and navigation.

In pursuing this range of participatory sensing applications we have noticed a unifying application style that we refer to as observational campaigns. The inspiration for this approach derives of course from well known techniques related to pub-sub and even earlier AI-developed techniques related to shared blackboards for coordination. In this context however, the approach gives context and focus to the sensory input, increasing our ability to treat the resulting data as sensory input that can be automatically processed. As an example, when there are flash floods (as are often experienced in LA), one could initiate a campaign to invoke the multiscale mobility that humans offer (both mobilizing to a location and then using local articulation to achieve standardized images) to collect images from specific distances and orientations and help focus civic resources on the most impacted of neighborhoods in semi-real time.

While we see tremendous public good driving adoption in the future, these sensing devices will be so personally and intimately placed that it just is not an option to leave privacy issues until later -- **neither in space nor time.** What I mean by that is that many people will never even turn on the device unless they can LOCALY filter the data…what we refer to as selective sharing. So this cant be left for 'later' either in terms of where or when that filtering or blurring is done. I find it fascinating that in this application context we have encountered another drier for in network, local processing…in this case for personal privacy. At the same time, as Steven Bellovin has pointed out, this will greatly compound issues associated with the human interface side of privacy and security mechanism configuration and verification.

And of course Privacy is highly context dependent and the nature of these policies and mechanisms will have to be conducive to the different contexts of public and private spaces.

The system components needed to support these intensively location and yet privacy sensitive applications are not new to Mobicom and extend far beyond the local devices to Internet elements that will support location verification and data aggregation and blurring. In our current relatively early stages of development, a lot of the key functionaility is located in what we refer to as trusted Mediator boxes that would sit at network edges and serve to attest to location and context information and also support data tagging and blurring as desired. Again, as with in network processing, we see some recurring themes carried over from our design of other types of wireless sensing systems such as multiscale sensing and actuation to achieve both density and coverage, and increased semantic support from network elements for namig, tagging, verifying, and blurring.

While things are quite early, and we will undoubtedly learn a lot about the specific architecture needs from more extensive development and use, it does seem safe to say that the Mobicom terrain of mobile, wireless systems, will have ever increasing impact as we exploit these devices as **embedded sensors**, and across a wide range of applications from civic to public health and natural resource management. And from the sensing perspective, this modality can be seen as filling in a critical gap in the family of sensing modalities that include remote sensing, static sensing, and automated mobility.

In conclusion, science applications have dominated the initial application of the embedding sensing systems described, however, the technology is expanding to monitoring agencies and dual use applications and its hard to imagine a segment of the economy, industrial enterprise, and daily life, that wont be affected.

Acknowledgments Work presented here is that of students, staff, and faculty at CENS with support from NSF #CCR-0120778, Nokia, Intel, Cisco, MSR.